arXiv:1609.08094v1 [physics.chem-ph] 26 Sep 2016

Efficient and portable acceleration of quantum chemical many-body methods in mixed floating point precision using OpenACC compiler directives

Janus J. Eriksen*

Institut für Physikalische Chemie, Johannes Gutenberg-Universität Mainz, D-55128 Mainz, Germany

E-mail: jeriksen@uni-mainz.de

Abstract

It is demonstrated how the non-proprietary OpenACC standard of compiler directives may be used to compactly and efficiently accelerate the rate-determining steps of two of the most routinely applied many-body methods of electronic structure theory, namely the second-order Møller-Plesset (MP2) model in its resolution-of-the-identity (RI) approximated form and the (T) triples correction to the coupled cluster singles and doubles model (CCSD(T)). By means of compute directives as well as the use of optimized device math libraries, the operations involved in the energy kernels have been ported to graphics processing unit (GPU) accelerators, and the associated data transfers correspondingly optimized to such a degree that the final implementations (using either double and/or single precision arithmetics) are capable of scaling to as large systems as allowed for by the capacity of the host central processing unit (CPU) main memory. The performance of the hybrid CPU/GPU implementations is assessed through calculations on test systems of alanine amino acid chains using one-electron basis sets

*To whom correspondence should be addressed

of increasing size (ranging from double- to pentuple-ζ quality). For all but the smallest problem sizes of the present study, the optimized accelerated codes (using a single multi-core CPU host node in conjunction with six GPUs) are found to be capable of reducing the total time-to-solution by at least an order of magnitude over optimized, OpenMP-threaded CPU-only reference implementations.

1 Introduction

When formulating the criteria for what defines *portable* source code, one tends to draw a distinction between two types of portability, namely that associated with functionality and that associated with performance. While the former of these covers what is perhaps usually implied when referring to code portability, i.e., the ability of a single code to run anywhere (rebuild the code on various architectures, the code will run and produce correct results), the latter, the existence of which have often been declared a fallacy,[1] relates to the ability of a single code to run productively anywhere, i.e., achieving, say, $80-90\%$ of the hand-tuned performance on any given architecture. As such, making a code portable will often result in a compromise between superior performance on a single platform and decent performance on all potential platforms. One possible way of circumventing this problem is by cluttering the source code with multiple conditional preprocessor directives and retain architecture-specific versions of key kernels with vendor-specific intrinsics. However, pursuing this strategy is bound to severely hamper two of the other main best practices of code development, namely those which are concerned with securing code *maintainability* and testability.[2]

For execution on standard multi-core central processing unit (CPU) architectures, the *OpenMP standard*[3]—first released in the second half of the 1990's as a means to parallelize loops across compute threads—has nowadays manifested itself as the de facto programming model for intranode worksharing and, in particular, single-instruction multiple-data (SIMD) parallel instructions. The advantages of employing OpenMP compiler directives are manifold; the standard is open and supported by most vendors, the compiler directives (or pragmas) are ignored entirely by those few compilers that do not offer support, and the syntax is relatively simple to comprehend, even to the uninitiated, making the learning barrier less steep than, e.g., that for intra- and internode message passing using the MPI protocol. Recently, with the advent of OpenMP versions 4.0 and 4.5, the standard now also allows for targeting other types of architectures such as many-core x86 processors, e.g., Intel® Xeon Phi™ products, and graphics processing units (GPUs), e.g., NVIDIA® or

AMD® GPUs. These feature enhancements from homo- to heterogenenous platforms allow for offloading of compute intensive code regions to coprocessors via so-called *device* constructs (in this terminology, the CPU, to which the coprocessor is connected, is denoted the *host*). However, there exist intrinsic difficulties in transferring the worksharing capabilities of OpenMP on the host to a corresponding worksharing on any attached device, in particular GPUs, as the architectural details of these are fundamentally different from those of the host. Specifically, the traditional OpenMP model of spawning a single team of threads, which then divide loop iterations among them, will not be suitable per se for execution on GPUs, as these are composed of multiple processing elements (PEs) that run in parallel (with no synchronization options between the threads in different PEs) and where each individual PE has the ability to efficiently perform vector-like operations. There are ways to bypass this issue, though, for instance by creating multiple teams of threads, across which the workload may be distributed. However, this is almost guaranteed to result in *accelerated* code that will not execute satisfactorily on the host (thereby losing performance portability), which will reintroduce the need for **(i)** preprocessing of different code blocks in addition to **(ii)** code duplication with the required maintenance that inevitably arises as a result of this.[4]

Instead, an alternative standard—the *OpenACC standard*,[5] recently founded by a consortium consisting of NVIDIA® and several key vendors, including Cray® and PGI®—has been established in an attempt to facilitate such portable and maintainable parallelism across various architectures, that is, provide the ability to expose to the compiler in a simple way what parts of the code can be run in parallel, and then leave it to the compiler to build the accelerated code *on a given platform*, be that the host (as a real alternative to OpenMP) or the device (as an inherently different approach to OpenMP). Since OpenACC is designed with heterogenous programming in mind, it is not biased towards multi-core CPU architectures in the same way as OpenMP is, while still being able to advantage from past endeavours made by compiler developers that tuned these towards vectorization. In particular, as opposed to adhering to a strictly prescriptive programming model

of simple thread teams, the OpenACC standard extends this by adding an additional layer of flexibility; for instance, when targeting GPUs, on which it is reasonable to think of a PE as a streaming multiprocessor (SM), the standard offers the possibility of parallelizing over threadblocks in a grid, i.e., individual SMs, warps (a group of 32 threads comprising the minimum execution unit), as well as individual CUDA threads within a warp. In the OpenACC terminology, these three layers are known as gangs, workers, and vectors.

The OpenACC programming model is a so-called host-directed sequential model capable of leveraging the parallel potential of one or even multiple accelerator devices, be that GPUs or multi-core coprocessors. The constructs, which the standard is designed to target, are the same as those the OpenMP standard are aimed at, namely nested loop structures, which themselves are suitable for shared memory multiprocessing. The translator between the input program and the individual SIMD units on the target architecture is the compiler, and the vectorization inhibitors are thus the same as for OpenMP, e.g., loop carried dependencies, prohibitive complexity (e.g., function calls), as well as indirect addressing and striding. As the compiler is directed by pragmas used to expose the parallel regions in the code, the code developer may approach the problem of offloading a certain kernel in a descriptive and, importantly, incremental manner, which makes for a less cumbersome workflow than if the entire code has to be implemented from scratch using some low-level approach unrelated to the language in which the actual code base is written (e.g., CUDA or OpenCL, in a GPU context). Furthermore, as the accelerated code will be based on the original source, this warrants, in turn, a functionally and performance portable, intuitively transparent, and easily extendable and maintainable implementation.

In the present work, it will be demonstrated how to make use of OpenACC compiler directives in the acceleration of electronic structure theories, with an illustrative application to two popular many-body methods, namely the second-order Møller-Plesset (MP2) model[6] in its resolution-of-the-identity (RI) approximated form and the preeminent perturbative (T) triples correction to the

coupled cluster[7] (CC) singles and doubles (CCSD) model,[8] compositely known as the CCSD(T) model.[9,10] To focus the discussion, the present work will be solely concerned with offloading to GPU accelerators, although we stress once again how the OpenACC standard, and hence the present approach, is completely general and, as such, not limited to nor biased against accelerators of a specific kind. On that note, we will start by briefly alluding to the existing literature, which—despite the use of general-purpose GPUs in natural sciences being a rather recent topic—is rich with work devoted to GPU-accelerated quantum chemistry, concerned with diverse topics ranging from the generation of electron repulsion integrals (ERIs),[11–15] over self-consistent Hartree-Fock (HF), complete active space (CAS), and density functional theory (DFT) methods,[16–24] to solvent models,[25,26] force fields,[27] and semi-empirical[28] as well as many-body methods.[29–38] Adding to this quantum, several program packages are today released with full or partial GPU-enabled features, e.g., Terachem,[39] NWChem,[40] among others.

That said, the adaption of GPUs and other types of coprocessors in the acceleration of quantum chemical and physical methods and applications as a whole has often been criticized for being too difficult a task to undertake implementation-wise and hence not worth the effort, in part also because of the uncertainty of (the shape of/general consensus on) future architectures that have always permeated the community (hesitation towards vectorization, massive parallelism, high-performance computing, etc.)[41] On top of that, it is important to note how the clear majority of all quantum chemical implementations in existence today have been developed and are being maintained by domain scientists, many of which hold temporary positions working (simultaneously) on various theoretical and/or application-oriented projects.[42] Thus, while said scientists might decide to implement a given accelerated method from scratch (again, presumably making use of some low-level approach), they might not be responsible for extending it with new features and capabilities in the future nor for the routine maintenance required with the emergence of novel architectures, not to mention the finer architectural diversities between different generations of, for instance, GPU accelerators. Add to that the typical requirements of a single code base and, in

particular, platform-independence, which most codes are subject to in the sense that any addition of accelerated code must not interfere with the standard compilation process (cf. the discussion in the opening paragraph of the present Section), and one will potentially arrive at more reasons against investing the necessary efforts than actual reasons in favor of doing so.

However, a recent paper from the Martínez group at Stanford University has outlined an alternative strategy for avoiding the above issues, albeit one targeted exclusively at GPU hardware.[43] The paper, which succeeds earlier work on so-called meta-programming from within the same group,[44] proposes the application of an automatic code generator capable of employing mathematical knowledge of the underlying algorithms (for the efficient generation of ERIs, in this case) to create a wide spectrum of possible code variants that explore, in their own words, the full space of all possible program implementations. This type of empirical, so-called profile-guided optimization subsequently tests the (multiple) candidate implementations on the target architecture at hand before deciding upon a final version, and the combination of a graph-based code generator, a tester, and an optimizer hence aims at producing the optimal implementation for previous and current generations of GPUs as well as, at least in principle, future generations.

These numerous code variants, which may be automatically produced using the engine in Ref. 43, will each make up different compromises between a number of parameters and factors; for instance, the efficiency of a given kernel will be tightly bound to its utilization of the complex memory hierarchy on the GPUs, the latency of which increases drastically upon traversing up through it, and the algorithmic complexity of the individual kernels trivially increases with respect to the maximum angular momentum of the one-electron basis set. Furthermore, the choice of resulting implementation will also depend on the choice of floating point precision, since lowering this from double to single precision—granted that the requested accuracy allows for it—may give preference to other variants. Importantly, however, the use of single precision arithmetics will result in an overall speed-up of the calculation, for execution on the host CPU and particularly

so for execution on GPUs. Now, in comparison with the generation of ERIs, the computation of correlation energies using (perturbative) many-body methods poses rather different requirements on the final implementation, as we will return to in Section 3. That being said, it is the intent of the present author to illustrate, through OpenACC implementations of the RI-MP2 and CCSD(T) methods, how simple and transparent such methods may be accelerated. It should be noted, however, that the adaption of an existing CPU ERI code to GPUs (or other types of accelerators, *vide supra*) by means of OpenACC directives will no doubt be more challenging, but it is still probably easier for the developer than using, e.g., CUDA and writing the code from scratch. In particular, the programmer may start by incrementally porting one kernel after another by exposing the inherent parallelism (loop structures) of the algorithm using directives, thereby verifying the correctness of the implementation along the way prior to subjecting the initial implementation to a subsequent optimization. In addition, the final implementation will require a minimum of maintenance, as the optimizer—with respect to novel architectures and compute capabilities—remains the compiler itself, and since the directives are treated as mere comments to a non-accelerating compiler, this applies for the efforts that will have to be invested into platform-independence of the source as well.

The present work is organized as follows. Following a brief introduction to the RI-MP2 and CCSD(T) methods in Section 2, the actual OpenACC implementations of these—using a combination of single and double precision arithmetics—are discussed in Section 3. In Section 4, we provide some computational details on the calculations to follow in Section 5, in which the performance of the implementations is assessed through calculations on test systems consisting of alanine amino acids in α-helix arrangements using one-electron basis sets of increasing size (ranging from double- to pentuple-ζ quality). Finally, a short summary alongside some conclusive remarks are presented in Section 6.

2 Theory

In electronic structure theory, CC and many-body perturbation theory[45] (MBPT) both offer a systematic approach towards the full configuration interaction[46] (FCI) wave function—the exact solution to the time-independent, non-relativistic electronic Schrödinger equation within a given one-electron basis set. In both hierarchies of methods, the mean-field HF solution acts as the reference to which correlated corrections are made by including excited configurations in the wave function. Moving beyond the HF approximation, not only Fermi correlation is considered, as also the correlated motion of electrons of opposite spin start being described. This implies that upon traversing up through either of the CC or MBPT hierarchies, increasingly more of such dynamical correlation is included in the wave function, with a complete description met at the target FCI limit.

In the MBPT series, the lowest-order perturbative correction to the HF energy for the isolated effect of connected double excitations is that of the non-iterative MP2 model.[6] As an improvement, the CC hierarchy offers the iterative CCSD model,[8] which accounts not only for connected double excitations, but does so to infinite order in the space of all single and double excitations out of the HF reference. However, in order to achieve relative energies of chemical accuracy ($\sim$ 1 kcal/mol), the CCSD energy has to be augmented by correction terms built from higher-level excitations. To a lowest approximation, corrections for effects due to connected triple excitations have to be considered. Among such higher-level approaches, the CCSD(T) model[9] is by far the most prominent. Unfortunately, the additional accuracy of the CCSD and, in particular, the CCSD(T) models over the MP2 model comes at a price, with the evaluation of the rate-determining step scaling non-iteratively as $\mathscr{O}(N^5)$ in the MP2 model (where N is a composite measure of the total system size), iteratively as $\mathscr{O}(N^6)$ in the CCSD model, and non-iteratively as $\mathscr{O}(N^7)$ for the (T) correction to the converged CCSD energy. Thus, whereas all three methods have nowadays become standard tools to computational quantum chemists, the current application range of the MP2 model considerably exceeds that of the CCSD(T) model, which are the two models that we will limit our focus to herein, due to the practical differences in associated computational cost.

2.1 RI-MP2

In spite of the fact that the MP2 model is formally far less expensive than the CCSD(T) model, the time-to-solution for the average MP2 calculation may still be substantial when targeting increasingly larger systems, e.g., in comparison with cheaper, albeit less rigorous and less systematic methods such as those of semi-empirical nature or those that calculate the energy from a functional of the one-electron density (DFT methods). For this reason, most efficient implementations of the MP2 model invoke an RI approximation for reducing the computational cost (the prefactor) as well as lowering the memory constraints.[47–50] Despite being an approximation to the MP2 model, the focused development of optimized auxiliary basis sets has managed to significantly reduce the inherent RI error, and for most reasonable choices of the intrinsic thresholds of a modern implementation, the error affiliated with the actual approximation will be negligible. However, while the computational cost is notably reduced in an RI-MP2 calculation with respect to its MP2 counterpart, its formal $\mathcal{O}(N^5)$ scaling with the size of the system will often still deem it demanding if no fundamental algorithmic changes are made.[51–55]

In the closed-shell (canonical) MP2 model, the correlation energy is given by the following expression

$$E_{\mathrm{MP2}} = -\sum_{ab}\sum_{ij} \frac{g_{aibj}(2g_{aibj} - g_{ajbi})}{\varepsilon_{ij}^{ab}} \tag{2.1.1}$$

in terms of two-electron ERIs, g_{aibj} (Mulliken notation), over spatial (spin-free) HF virtual orbitals a, b and occupied orbitals i, j, as well as the difference in energy between these

$$\varepsilon_{ij}^{ab} = \varepsilon_a + \varepsilon_b - (\varepsilon_i + \varepsilon_j) \; . \tag{2.1.2}$$

Besides the final evaluation of the MP2 energy in Eq. (2.1.1), the dominant part of an MP2 calculation is the construction of the four-index ERIs. In the RI-MP2 method, these are approximated

by products of two-, $V_{\alpha\gamma}$, and three-index, $W_{ai\alpha}$, integrals by means of the following symmetric decomposition

$$g_{aibj} \approx \sum_{\gamma} C_{ai}^{\gamma} C_{bj}^{\gamma} \, . \tag{2.1.3}$$

In Eq. (2.1.3), greek indices denote atomic orbitals within an auxiliary fitting basis set used for spanning the RI, and the fitting coefficients, C_{ai}^{γ}, are defined as

$$C_{ai}^{\gamma} = \sum_{\alpha} W_{ai\alpha} [\mathbf{V}^{-1/2}]_{\alpha\gamma} \, . \tag{2.1.4}$$

In terms of computational cost, the evaluation of two- and three-index integrals and/or the calculation of the fitting coefficients will dominate the overall calculation for small systems ($\mathscr{O}(N^3)$- and $\mathscr{O}(N^4)$-scaling processes, respectively). However, upon an increase in system size, the final $\mathscr{O}(N^5)$-scaling assembly of the two-electron integrals in Eq. (2.1.3) will start to dominate, and this process is thus an ideal candidate for accelerating the RI-MP2 model, cf. Section 3.

2.2 CCSD(T)

As touched upon in the opening paragraph of the present Section, the (T) correction scales non-iterative as $\mathscr{O}(N^7)$, as opposed to the $\mathscr{O}(N^6)$-scaling iterative evaluation of the energy and cluster amplitudes in the preceding CCSD calculation. Thus, for all but the smallest systems, the evaluation of the (T) correction, at least in its conventional canonical formulation, will dominate. In a basis of local HF orbitals, however, this balance might shift, making the underlying evaluation of the CCSD energy (and amplitudes) the computational bottleneck.[56–63]

For closed-shell systems, the CCSD(T) energy is defined by a triples correction, $E_{(\mathrm{T})}$, to the CCSD energy, in turn defined by two energy contributions rationalized from MBPT: $E^{[4]}$, a fourth-order term involving CCSD doubles amplitudes, $\{t_{ij}^{ab}\}$, and $E^{[5]}$, a fifth-order term involving CCSD

singles amplitudes, $\{t_i^a\}$. In a basis of canonical HF spatial orbitals, these may be expressed as[64]

$$E^{[4]} = \sum_{abc}\sum_{ijk} \tilde{t}_{ijk}^{abc} t_{ijk}^{abc} \varepsilon_{ijk}^{abc} \tag{2.2.1a}$$

$$E^{[5]} = \sum_{abc}\sum_{ijk} \tilde{z}_{ijk}^{abc} t_{ijk}^{abc} \varepsilon_{ijk}^{abc} \ . \tag{2.2.1b}$$

In Eq. (2.2.1), ε_{ijk}^{abc} is defined on par with ε_{ij}^{ab} in Eq. (2.1.2), and the triples amplitudes, $\{t_{ijk}^{abc}\}$, are defined as

$$t_{ijk}^{abc} = -P_{ijk}^{abc} \frac{\sum_d t_{ij}^{ad} g_{ckbd} - \sum_l t_{il}^{ab} g_{ckl\,j}}{\varepsilon_{ijk}^{abc}} \tag{2.2.2}$$

where P_{ijk}^{abc} is a symmetrization operator

$$P_{ijk}^{abc} x_{ijk}^{abc} = x_{ijk}^{abc} + x_{ikj}^{acb} + x_{jik}^{bac} + x_{jki}^{bca} + x_{kij}^{cab} + x_{kji}^{cba} \ . \tag{2.2.3}$$

The z_{ijk}^{abc} coefficients in Eq. (2.2.1b) are given as

$$z_{ijk}^{abc} = -\frac{(t_i^a g_{jbkc} + t_j^b g_{iakc} + t_k^c g_{iajb})}{\varepsilon_{ijk}^{abc}} \tag{2.2.4}$$

and an arbitrary six-index quantity $\tilde{x}_{ijk}^{abc}$ in Eq. (2.2.1) is defined as

$$\tilde{x}_{ijk}^{abc} = 2(\tfrac{2}{3} x_{ijk}^{abc} - x_{ijk}^{acb} + \tfrac{1}{3} x_{ijk}^{bca}) \ . \tag{2.2.5}$$

The rate-determining step is now identified as being the construction of the triples amplitudes in Eq. (2.2.2), whereas the final evaluation of $E_{(\mathrm{T})}$ only scales as $\mathscr{O}(N^6)$. As opposed to the evaluation of the RI-MP2 energy in Section 2.1, for which only parts of the algorithm will be accelerated (Eq. (2.1.1) and Eq. (2.1.3) combined), all of the (T) kernels will be offloaded to the accelerator(s), as will be detailed in the following Section 3.

3 Implementations

As is obvious from Section 2, the evaluation of both the RI-MP2 energy and the (T) correction relies on a vast number of matrix-matrix multiplications (MMMs), which in all modern implementations of the methods are handled by making calls to the `dgemm` routine of an optimized level-3 basic linear algebra subprograms[65] (BLAS-3) library. The fact that both methods may be recasted such that they utilize BLAS libraries optimized for modern multi-core CPU architectures (e.g., Intel® MKL, ATLAS,[66] or GotoBLAS,[67] to name a few) implies that they are amenable to accelerators as well (GPUs, in the present context), for which optimized BLAS libraries too exist, e.g., MAGMA[68] and CUBLAS.[69] Thus, in contrast to the generation of ERIs, cf. the discussion in Section 1, which is comprised of extensive parallel regions, each with numerous tightly nested loop structures, this forces somewhat different requirements on an accelerated implementation for coprocessors; in particular, aspects such as efficient data movement patterns and data reuse become highly important, since the tensors involved in many-body methods tend to grow large (in excess of what might be afforded by the available device memory) and the current interconnects between a host and its attached accelerator(s) (e.g., a PCIe bus between a CPU host and a GPU) generally impede unnecessary (or redundant) movement of data. The few remaining loop constructs (*vide infra*) are then translated by descriptive directives into kernels that can run in parallel on the accelerator. By letting the host processor orchestrate the execution of all parallel kernels, those cores on the host that do not control an accelerator are free to participate in the workload. This is an important feature; by allowing for the final code to execute in a hybrid manner on both the host CPU as well as the attached accelerator device(s), the present approach is not biased towards neither the former nor the latter of the two platforms (making for a fair comparison of the relative CPU vs. GPU throughput[70]). If the implementation was instead formulated in terms of some low-level language such as CUDA, such a bias would be present. In addition, the OpenACC standard further provides features for circumventing the default host-directed behavior, in which a sequence of operations gets scheduled for sequential execution on the device, by instead allowing for these to overlap in different pipelines, i.e., executing operations (computations and data transfers) asynchronously,

given that all potential race conditions have been appropriately mapped out.

Throughout the years, elaborate attempts have been made at pushing the limits for the RI-MP2[71–74] and CCSD(T)[75–81] methods by devising massively parallel implementations of the standard canonical formulations in Section 2. In addition, the methods have previously been ported to GPUs within the groups of Alán Aspuru-Guzik[29] (RI-MP2), Nakajima[30] (RI-MP2), Mark Gordon[34] (CCSD(T)), and Karol Kowalski[37] (CCSD(T)), as, e.g., reviewed in recent contributions to a book devoted to the subject of electronic structure calculations on graphics units.[82,83] In most, if not all of these implementations, the main driver routine, regardless of the method, consists of an outer nested loops, most often over occupied orbital indices with restricted summations to incorporate permutational symmetry of the two-electron ERIs.[84,85] For both methods, each loop cycle then involves (numerous) large MMMs, possibly intertwined by tensor (index) permutations and the division by orbital energy differences. Since each loop cycle is data-independent of all others, these may be trivially distributed over multiple CPU nodes using MPI.

In a likeminded fashion, OpenMP worksharing directives may be used to parallelize the individual loop cycles over the available (N_{threads}, or more generally, `OMP_NUM_THREADS`) cores on the host CPU node. In this case, N_{threads} single-threaded MMMs are executed concurrently, which will generally be more efficient than running a single multi-threaded MMM (over N_{threads} threads) sequentially N_{threads} times, albeit at the expense of having to store multiple private copies of some of the involved intermediates. As a prerequisite for using OpenMP this way, one needs to flatten (or collapse) the two (in the case of the RI-MP2 method, cf. Eq. (2.1.1)) or three (in the case of the (T) correction, cf. Eq. (2.2.1)) outer loops into one large, composite outer loop with loop cycles of identical weight. This way, the evaluation of the RI-MP2 energy or (T) correction is programmed to execute in a hybrid heterogeneous manner on both the CPU (host) cores and one or even multiple accelerator devices; combined, the accelerators may be regarded as augmenting the host CPU by additional cores (one for each attached accelerator), albeit some that operate at a

significantly higher clock rate. As far as the use of the GPUs is concerned, all of the required data blocks (tiles of the fitting coefficient tensor in the RI-MP2 code, and tiles of the involved CCSD doubles amplitudes and ERIs in the CCSD(T) code) can be transferred asynchronously, thereby overlapping all data movement between the host and the device(s) with computations on the latter, and for the actual tensors, we allocate these directly into page-locked (or pinned) memory in order to lower the cost of the data transfers.

Recently, in an invited chapter to a new book devoted to parallel programming with OpenACC, the present author outlined in some detail how to incrementally accelerate an OpenMP-parallelized CPU-only version of the RI-MP2 $\mathcal{O}(N^5)$-scaling kernel (Eq. (2.1.1) and Eq. (2.1.3) combined) by means of OpenACC compiler directives.[86] Herein, the (T) correction in Section 2.2 is accelerated in a similar fashion, and the source code for both implementations (RI-MP2 and the (T) correction) accompany the present work under an MIT license.[87] Furthermore, we will here assess the potential of lowering the numerical precision in certain parts of the RI-MP2 and (T) kernels. In standard implementations of the two methods, all of the involved steps (evaluation of ERIs, amplitudes, fitting coefficients, etc., as well as the final energy assemblies) are performed using double precision arithmetics exclusively. However, motivated by two recent papers,[88,89] which claim that single precision arithmetics can be applied in the construction of the four-index ERIs (Eq. (2.1.3)) used for RI-MP2 as well as in the construction of the triples amplitudes (Eq. (2.2.2)) for CCSD(T), we will do exactly this by substituting selective calls to `dgemm` with corresponding calls to `sgemm` in the codes. Not only is the performance bound to improve from the use of single precision arithmetics, but also the storage requirements on the GPU main memory will be lowered.

4 Computational details

For the purpose of evaluating the implementations of the two methods in double or mixed floating point precision, we will conduct performance tests on chains (α-helix arrangements) of ala-

nine amino acids ([ala]-n, where n is the number of residues) in cc-pVXZ (X = D, T, Q, and 5) one-electron basis sets[90] and corresponding cc-pVXZ-RI auxiliary basis sets[91,92] (in the case of RI-MP2). Due to the difference in application range, the RI-MP2 and CCSD(T) implementations will be tested for different problems sizes (a composite measure of the number of electrons and number of basis functions); specifically, the RI-MP2 implementations are tested for the [ala]-6 system in the cc-pVXZ (X = T, Q, and 5) basis sets as well as the systems [ala]-7 to [ala]-10 in a cc-pVQZ basis, whereas the CCSD(T) implementations are tested for the [ala]-1 system in the cc-pVXZ (X = D, T, and Q) basis sets as well as the systems [ala]-2 to [ala]-6 in a cc-pVDZ basis.

In terms of computational hardware, the accelerators used are NVIDIA® Kepler K40 GPUs (2880 processor cores @ 745 MHz (GPU Boost @ 875 MHz enabled for all calculations) and 12 GB main memory, from here on simply abbreviated as 'K40s') and the host nodes are Intel® Ivy Bridge E5-2690 v2, dual socket 10-core CPUs (20 cores @ 3.00 GHz and 128 GB main memory), i.e., `OMP_NUM_THREADS` = $N_{\text{threads}} = 20$ if not otherwise noted. The host math library is Intel® MKL (version 11.2) and the corresponding device math library is CUBLAS (CUDA 7.5). All calculations are serial (non-MPI), and the OpenMP-/OpenACC-compatible Fortran compiler used is that of the PGI compiler suite (version 16.4).

In all implementations, the tensors containing orbital energies, fitting coefficients, CCSD singles and doubles amplitudes, as well as ERIs are initialized with unique, yet arbitrary single and/or double precision numbers (depending on the context), as the calculation of these is not the objective of the present study, cf. Ref. 86. The Git hash of the code in Ref. 87 used for the actual production runs is `42f76337`.

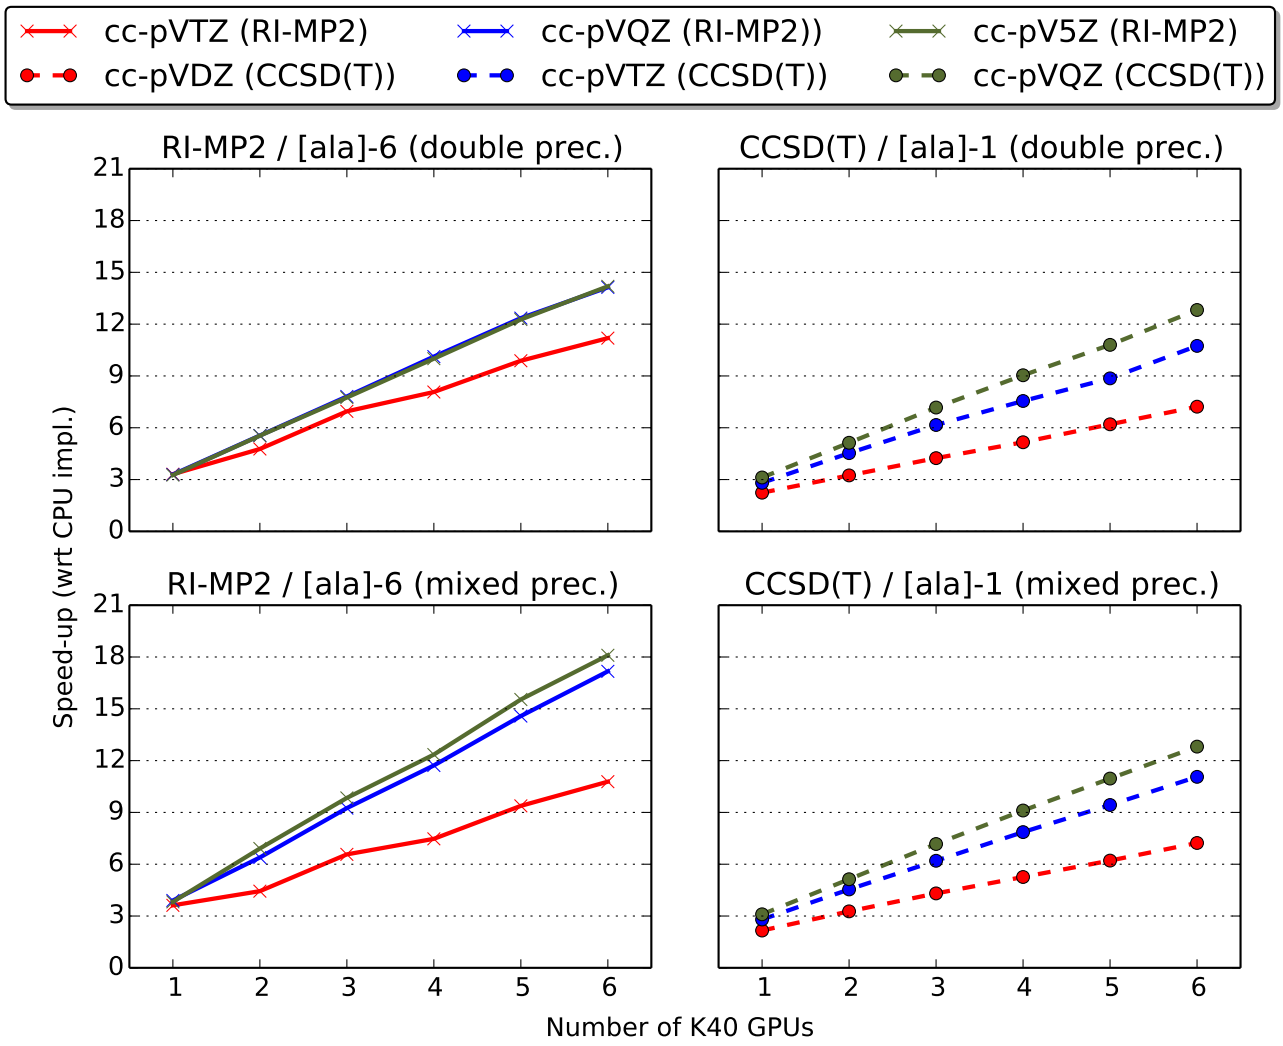

Figure 1: Multi-GPU scaling with the size of the basis set for the RI-MP2 and CCSD(T) implementations using double or mixed precision arithmetics.

5 Results

In Figure 1 and Figure 2, results for the total speed-up over the threaded CPU-only implementations (using a total of 20 OpenMP threads) are reported for fixed system size–varying basis sets (Figure 1) and varying system sizes–fixed basis set (Figure 2) as a function of the number of K40s involved in the calculations. Focusing first on the double precision RI-MP2 results, we note from Figure 1 how the [ala]-6/cc-pVTZ problem (by far the smallest of all the systems in the present RI-MP2 test) is too small for the use of the K40s to be optimal, while for all other combinations of system sizes and basis sets in Figure 1 and Figure 2, the use indeed appears so (*vide infra*). From

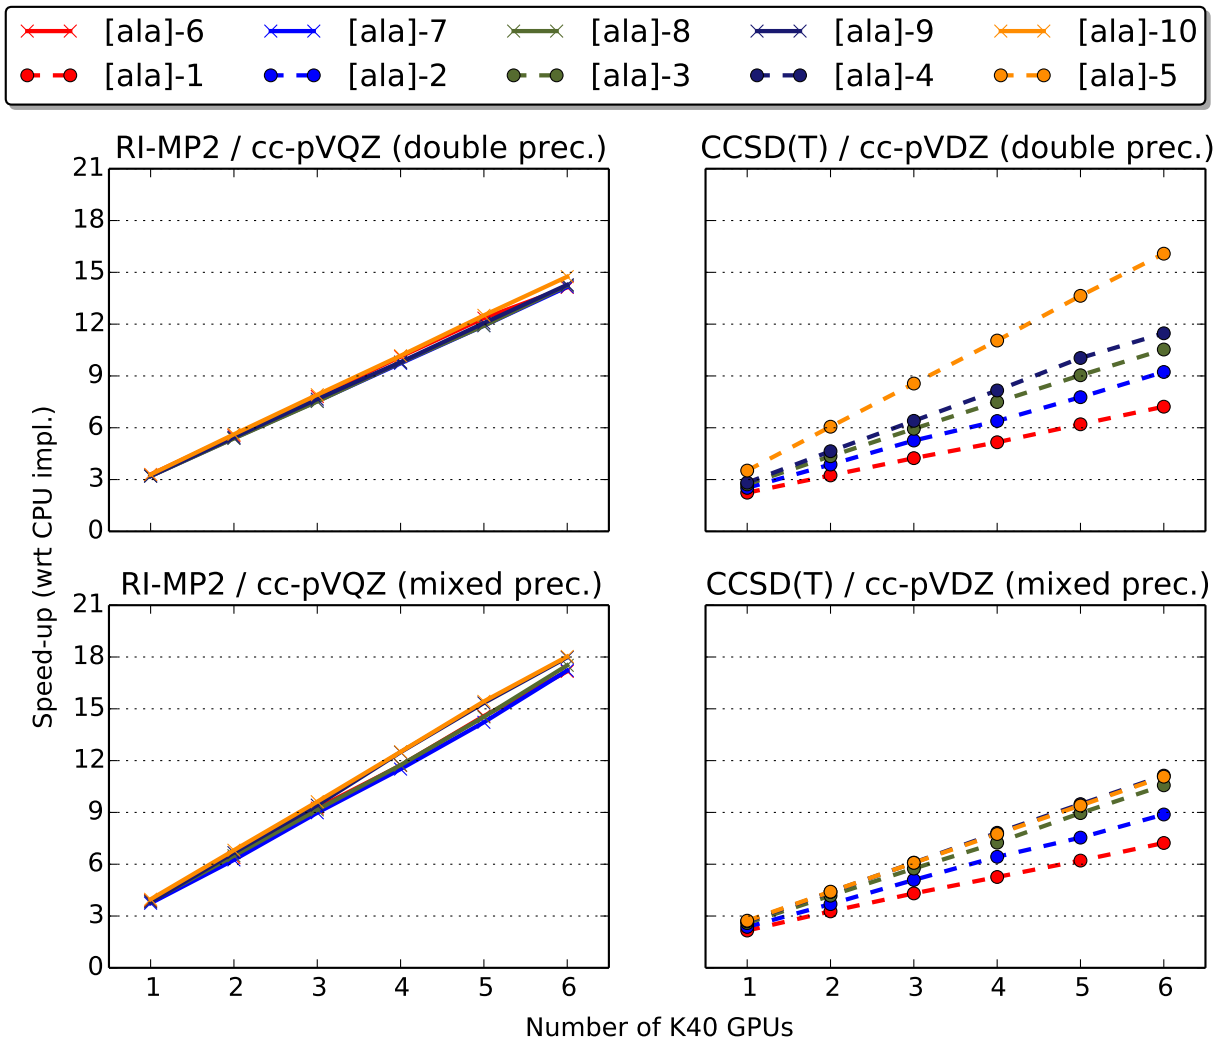

Figure 2: Multi-GPU scaling with the size of the molecular system for the RI-MP2 and CCSD(T) implementations using double or mixed precision arithmetics.

the double precision CCSD(T) results, on the other hand, a performance improvement is observed in both cases, with no obvious sign of saturation visible neither with respect to the size of the employed basis set in Figure 1 nor with respect to system size in Figure 2. Moving from double to mixed precision arithmetics, the multi-GPU scaling of the RI-MP2 implementation is observed to be unaffected for the [ala]-6/cc-pVTZ problem, but significantly improved for all other tests; in particular, the performance is even improved in the transition from a quadruple- to a pentuple-ζ basis set in Figure 1. For the mixed precision CCSD(T) implementation, however, the picture is much the same as for the equivalent implementation in double precision, although the improvement in

scaling with increasing system size is now not as prominent. In explaining why this is so, we recall how only parts of the overall calculation of the (T) correction is performed using single precision arithmetics (the construction of the triples amplitudes), as opposed to the RI-MP2 implementation, in which only the final (computationally insignificant) energy assembly is performed in double precision numbers. Thus, the speed-up from calling `sgemm` over `dgemm` will be smaller. Also, the dimensions of the involved MMMs for the tested CCSD(T) problem sizes are greatly reduced with respect to those of the RI-MP2 problem sizes—a practical necessity, due to the larger memory and cost requirements of the former of the two methods—which further impedes (or disfavors) the scaling of the single precision implementation in the case of CCSD(T). The use of single precision does, however, allow for even larger problem sizes to be addressed, as we will see later on, since the overall CPU and GPU memory requirements are reduced.

Having assessed the accumulated speed-up over the ordinary CPU-only implementations, we next report results for the actual scaling with the number of K40s. In Figure 3 and Figure 4, these scalings are presented, with the relative deviation from ideal behavior written out in the limit of six GPUs. As opposed to if the calculations were executed exclusively on the GPUs, i.e., in the non-hybrid limit where $N_{\text{threads}} == N_{\text{GPUs}}$, for which the ideal scaling is the trivial proportionality with the number of GPUs, N_{GPUs} (performance doubling, tripling, etc., on two, three, etc., GPUs), this is not the case for hybrid CPU/GPU execution, i.e., whenever $N_{\text{threads}} > N_{\text{GPUs}}$, as each CPU core is now treated as an accelerator on its own. Thus, the ideal speed-up for the latter, heterogeneous case is defined as

$$R = \frac{(N_{\text{threads}} - N_{\text{GPUs}}) + N_{\text{GPUs}} S}{(N_{\text{threads}} - 1) + S} \,. \tag{5.0.6}$$

In the definition of R in Eq. (5.0.6), the constant factor $S = N_{\text{threads}} K$, where K is the time ratio between a CPU-only calculation ($N_{\text{threads}} =$ `OMP_NUM_THREADS`; $N_{\text{GPUs}} = 0$) and a GPU-only calculation using a single GPU ($N_{\text{threads}} = N_{\text{GPUs}} = 1$), accounts for the relative difference in processing power between a single CPU core (assuming ideal OpenMP parallelization on the host)

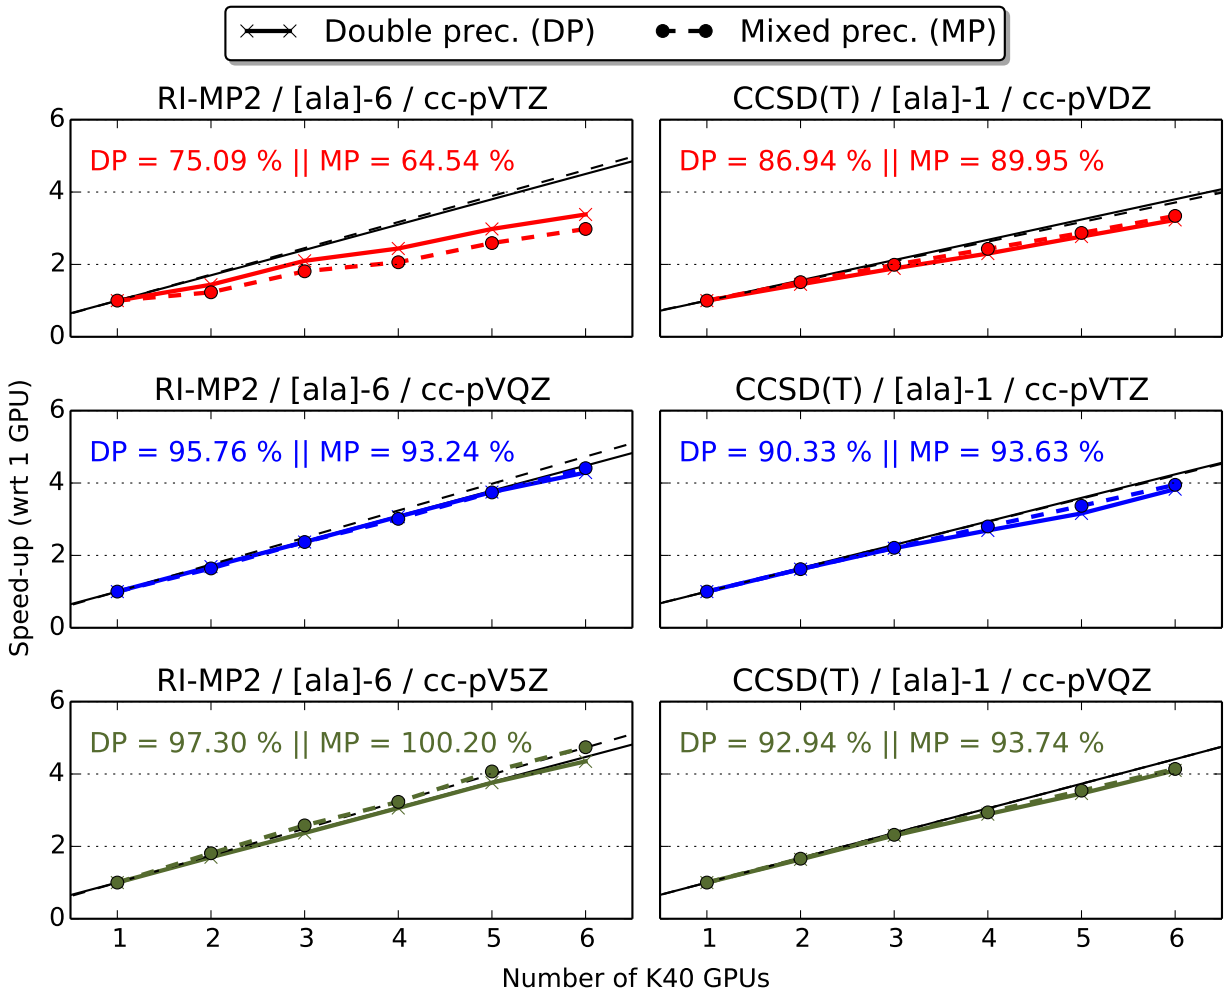
Double prec. (DP)
Mixed prec. (MP)
RI-MP2 / [ala]-6 / cc-pVTZ
DP = 75.09 % || MP = 64.54 %
CCSD(T) / [ala]-1 / cc-pVDZ
DP = 86.94 % || MP = 89.95 %
RI-MP2 / [ala]-6 / cc-pVQZ
DP = 95.76 % || MP = 93.24 %
CCSD(T) / [ala]-1 / cc-pVTZ
DP = 90.33 % || MP = 93.63 %
RI-MP2 / [ala]-6 / cc-pV5Z
DP = 97.30 % || MP = 100.20 %
CCSD(T) / [ala]-1 / cc-pVQZ
DP = 92.94 % || MP = 93.74 %
Speed-up (wrt 1 GPU)
Number of K40 GPUs

Figure 3: Deviation from ideal scaling with the number of K40s for combinations of a fixed system size (RI-MP2/[ala]-6 and CCSD(T)/[ala]-1) and varying basis sets (cc-pVXZ where X = D, T, Q, and 5). For details, please see the text.

and a single GPU. From the results in Figure 3 and Figure 4, we note how the scaling becomes (near-)ideal for all but the smallest problem sizes, a statement valid for both models in both double and mixed precision, regardless of the fact that these dimensions are significantly reduced in the CCSD(T) tests, cf. the discussion above.

In addition, we may monitor how large a percentage of the actual computations is being handled by the GPUs in the hybrid RI-MP2 and CCSD(T) implementations by noting how these involve a

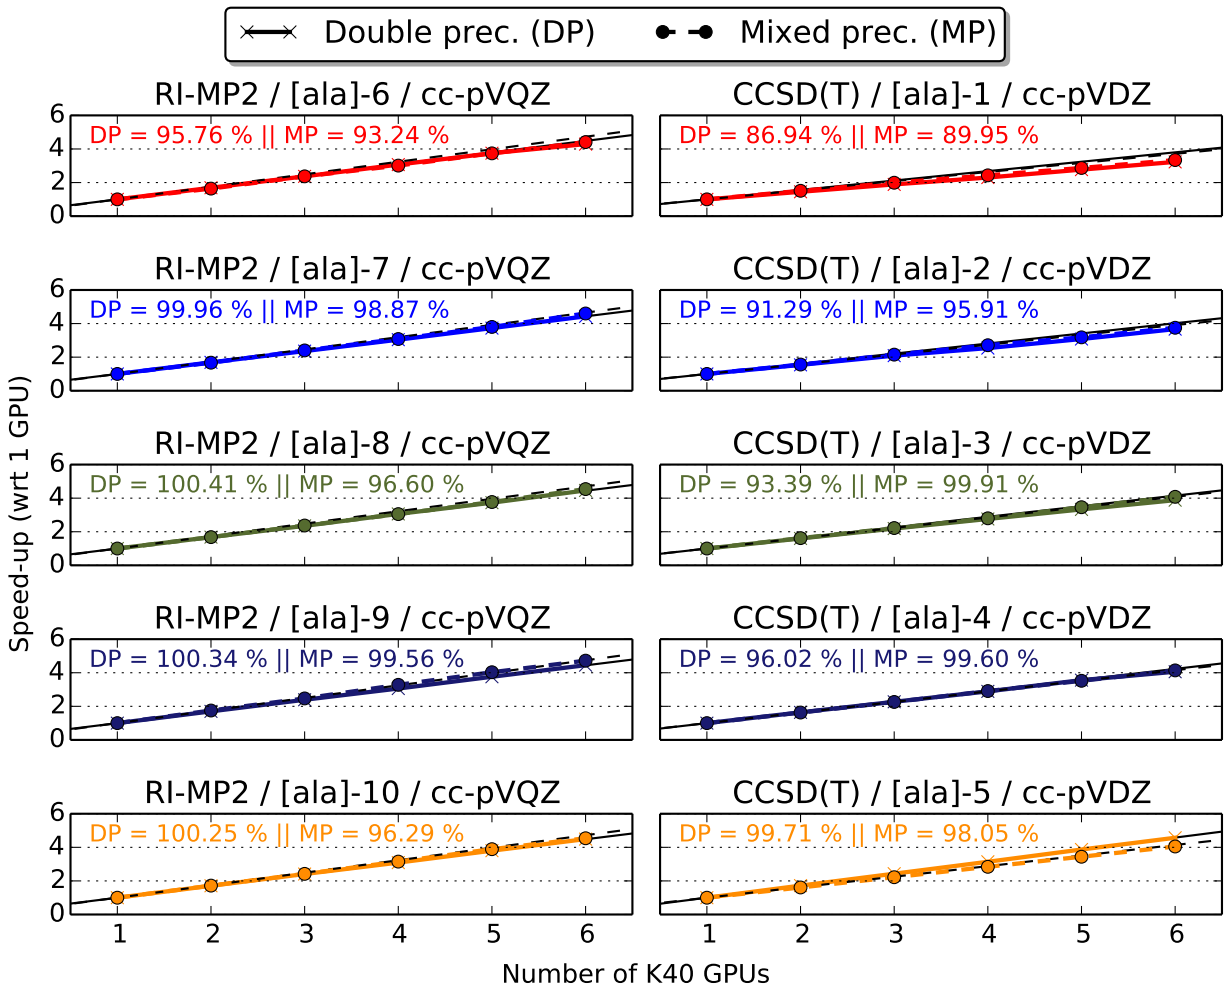
Double prec. (DP)
Mixed prec. (MP)
RI-MP2 / [ala]-6 / cc-pVQZ
DP = 95.76 % || MP = 93.24 %
CCSD(T) / [ala]-1 / cc-pVDZ
DP = 86.94 % || MP = 89.95 %
RI-MP2 / [ala]-7 / cc-pVQZ
DP = 99.96 % || MP = 98.87 %
CCSD(T) / [ala]-2 / cc-pVDZ
DP = 91.29 % || MP = 95.91 %
RI-MP2 / [ala]-8 / cc-pVQZ
DP = 100.41 % || MP = 96.60 %
CCSD(T) / [ala]-3 / cc-pVDZ
DP = 93.39 % || MP = 99.91 %
RI-MP2 / [ala]-9 / cc-pVQZ
DP = 100.34 % || MP = 99.56 %
CCSD(T) / [ala]-4 / cc-pVDZ
DP = 96.02 % || MP = 99.60 %
RI-MP2 / [ala]-10 / cc-pVQZ
DP = 100.25 % || MP = 96.29 %
CCSD(T) / [ala]-5 / cc-pVDZ
DP = 99.71 % || MP = 98.05 %
Speed-up (wrt 1 GPU)
Number of K40 GPUs

Figure 4: Deviation from ideal scaling with the number of K40s for combinations of varying system sizes ([ala]-n where $n = 1-10$) and a fixed basis set (RI-MP2/cc-pVQZ and CCSD(T)/cc-pVDZ). For details, please see the text.

total of $T_{\text{RI-MP2}}$ and $T_{\text{CCSD(T)}}$ tasks of equal weight

$$T_{\text{RI-MP2}} = N_{\text{o}}(1+N_{\text{o}})/2 \tag{5.0.7a}$$

$$T_{\text{CCSD(T)}} = \sum_{i=0}^{N_{\text{o}}-1} (N_{\text{o}}-i)(1+(N_{\text{o}}-i))/2 \tag{5.0.7b}$$

where N_{o} denotes the number of occupied orbitals. Through a dynamic OpenMP schedule, these individual tasks are distributed among the CPU cores and K40s in either of the implementations, so the scaling results in Figure 3 and Figure 4 may be complemented by corresponding results for the relative GPU workload. From the results, presented in Figure 5 as the accumulated workload in

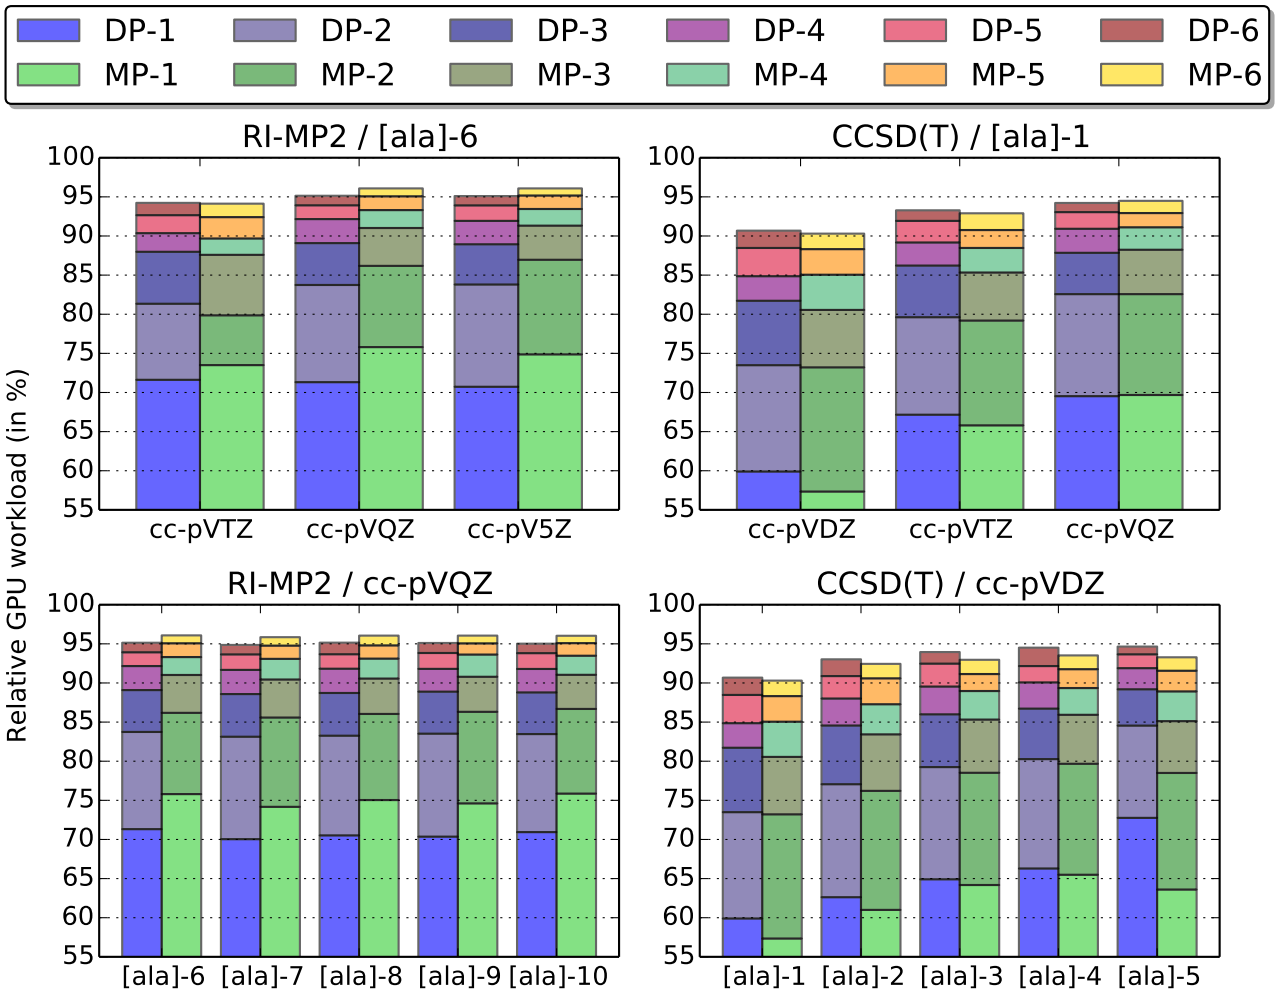

Figure 5: Accumulated relative GPU workload with number of K40s (DP-n/MP-n where $n = 1-6$) using the implementations in either double (DP) or mixed precision (MP) numbers.

percent, we note how for the present problem sizes and CPU/GPU hardware, the actual utilization of the host node is minor (less than 10%) when, say, three or more GPUs are attached to said node, regardless of the model. Still, the hybrid schemes are always faster than non-hybrid analogues, i.e., when $N_{\text{threads}} == N_{\text{GPUs}}$.

Finally, we compare the total time-to-solution for the CPU-only and hybrid CPU/GPU implementations of the RI-MP2 and CCSD(T) models in Figure 6.[93] From these results, using six K40s, as in the present study, is seen to reduce the total time-to-solution over the CPU-only implementations—in either double or mixed precision—by at least an order of magnitude for all but the smallest possible problem sizes. This is indeed a noteworthy acceleration of both models.

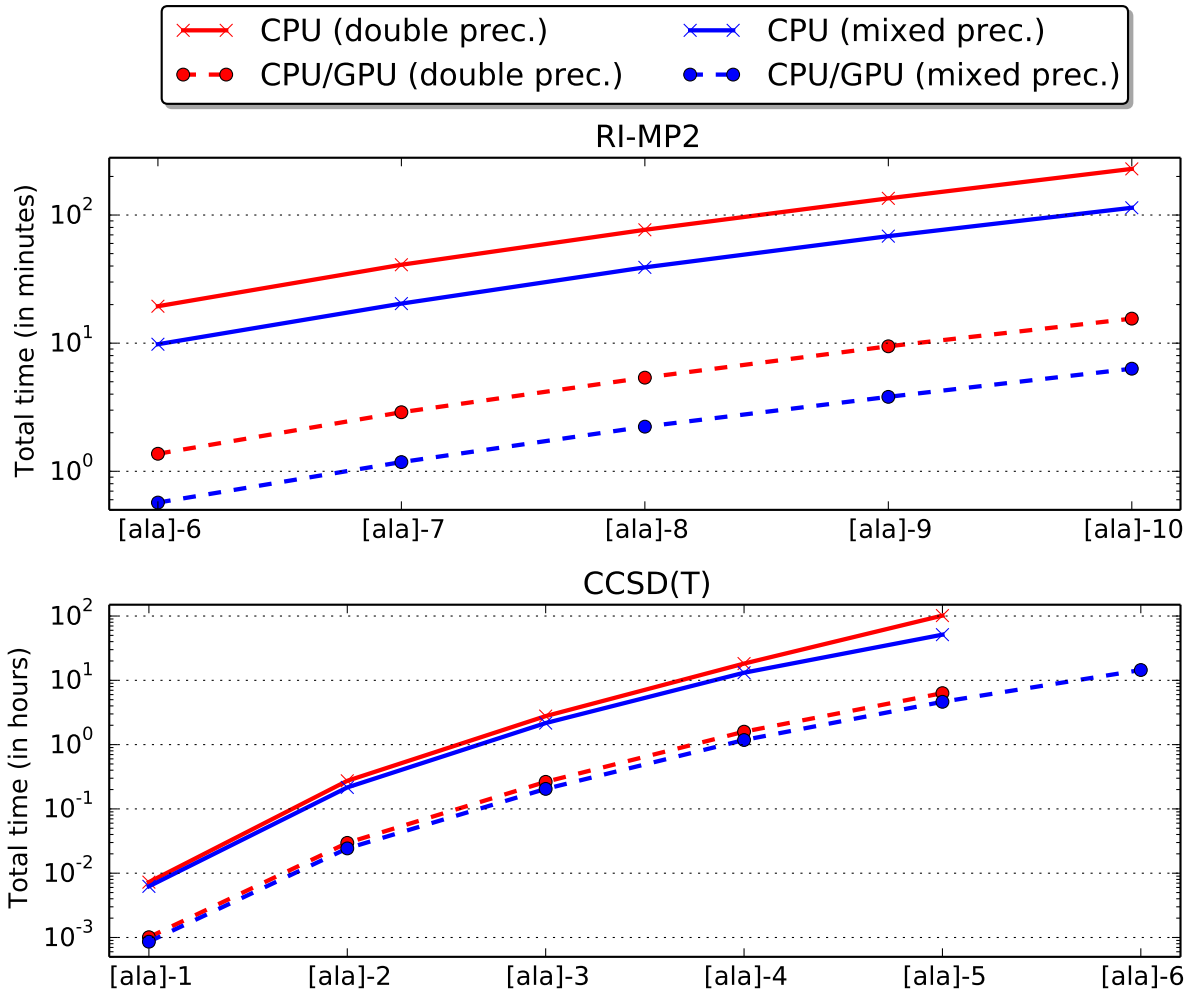

Figure 6: Total time-to-solution for the CPU-only and hybrid CPU/GPU implementations (using six K40s) of the RI-MP2 and CCSD(T) models.

In particular, we note how for the RI-MP2 model, the accelerated calculation on the [ala]-10 system in a cc-pVQZ basis ($N_\mathrm{o} = 195$, $N_\mathrm{v} = 4170$, and $N_\mathrm{aux} = 9592$ where N_v and N_aux are the number of virtual orbitals and auxiliary basis functions, respectively) took less time than the corresponding CPU-only calculation on the significantly smaller [ala]-6 system within the same basis ($N_\mathrm{o} = 119$, $N_\mathrm{v} = 2546$, and $N_\mathrm{aux} = 5852$), while for the calculation of the (T) correction, an accelerated calculation on the [ala]-6/cc-pVDZ system ($N_\mathrm{o} = 119$ and $N_\mathrm{v} = 475$, cf. the note in Ref. 93) terminates in less time than a corresponding CPU-only calculation on the [ala]-4 system within the same basis ($N_\mathrm{o} = 81$ and $N_\mathrm{v} = 323$). On the basis of these results, it may be argued that the use of a combination of OpenMP and OpenACC compiler directives—as long as the complete fitting coefficients,

integrals, CCSD amplitudes, etc., fit into main memory on the host—makes it somewhat unnecessary to explicitly parallelize the code using MPI with the complications that inevitably arise from this.

6 Summary and conclusion

In the present work, OpenACC compiler directives have been used to compactly and efficiently accelerate the $\mathcal{O}(N^5)$- and $\mathcal{O}(N^7)$-scaling rate-determining steps of the evaluation of the RI-MP2 energy and (T) triples correction, respectively. In the accelerated implementations, the kernels have all been offloaded to GPUs, and an asynchronous pipelining model has been introduced for the involved computations and data traffic. Due to their minimal memory footprints and efficient dependence on optimized math libraries, the implementations using either double or a combination of both single and double precision arithmetics are practically capable of scaling to as large systems as allowed for by the capacity of the host main memory.

In Ref. 43, it was argued that one cannot completely rely on a general-purpose compiler in the search for better program transformations since such an imaginary object as an ideal compiler cannot exist, with reference to the somewhat demoralizing full employment theorem for compiler writers which states that for any optimizing compiler there will always exist a superior.[94] While this is perfectly true, instead of viewing it as a motivation for attempting to *beat the compiler* by creating additional compilation layers, such as code generators, testers, etc., one might take an alternative stand by obeying to a sort of ad hoc, yet generalizable philosophy of *embracing the compiler*; namely, given that an optimal coprocessor has not yet been mutually decided upon (for obvious reasons, as different vendors seek to promote their own variant), and given that consensus is yet to be reached on whether or not there is indeed a need for accelerators in the field of electronic structure theory, one might instead try to make the most out of what hardware is currently available (as well as near-future prospective hardware) by investing only the least possible amount

of effort into porting one's code to this. The present work is intended to argue the case that the compiler directives of the OpenACC standard serve exactly this purpose by providing the means for targeting various coprocessors (e.g., GPUs or many-core x86 processors) in addition to the multi-core host node itself in an *efficient*, *transparent*, and *portable* high-level manner.

While the performance degradation of an OpenACC implementation (with respect to a hand-tuned low-level implementation) is bound to be larger for more complex electronic structure algorithms, such as the generation of ERIs or mean-field HF and DFT methods, than in the present case of MBPT methods, the application to the evaluation of the RI-MP2 energy and (T) triples correction is intended to illustrate a number of key advantages of the use of compiler directives over a reformulation of an optimized implementation for CPUs in terms of, e.g., CUDA or OpenCL for execution on GPUs. First and foremost, it is the opinion of the present author that accelerated code needs to be relatively easy and fast to implement, as new bottlenecks are bound to appear as soon as one part of a complex algorithm has been ported to accelerators (cf. Amdahl's law[95]). Second, the use of compiler directives guarantees—on par with the use of OpenMP worksharing directives for SIMD instructions on standard CPU architectures—that the final code remains functionally portable, i.e., the addition of accelerated code does not interfere with the standard compilation of the code on commodity hardware using standard non-accelerating compilers. Third, since the RI-MP2 and CCSD(T) methods alongside other, more advanced non-iterative CC many-body methods alike[96–100] are intrinsically reliant on large matrix-vector and matrix-matrix operations, the main need for accelerators in this context is for offloading exactly these. Thus, besides a number of small generic kernels, e.g., tensor index permutations or energy summations, compiler directives are primarily used for optimizing the data transfers between the host and the device(s), for instance by overlapping these with device computations. Hopefully, the generality of the discussion in the present work will encourage and possibly aid others to accelerate similar codes of their own. As a means to facilitate exactly this, the present implementations come distributed alongside this work for others to reuse in part or even in full.[87]

Finally, one particular potential area of application for the present implementations deserves a dedicated mentioning. While the discussion of the methods herein has been exclusively concerned with their standard canonical formulations for full molecular systems, we note how both methods have also been formulated within a number of so-called local correlation schemes, of which one branch relies on either a physical[101,102] or orbital-based[103–105] fragmentation of the molecular system. In these schemes, standard (pseudo-)canonical calculations are performed for each of the fragments before the energy for the full system is assembled at the end of the total calculation. Thus, by accelerating each of the individual fragment (and possible pair fragment) calculations, the total calculation will be accelerated as well without the need for investing additional efforts, and the resulting reduction in time-to-solution hence has the potential to help increase the range of application even further for these various schemes.

Acknowledgments

The author gratefully acknowledges support from the NVIDIA® Corporation, in particular for granting access to an internal test cluster, the resources of which were used in the preparation of the present work. The author furthermore wishes to thank Mark Berger, Jeff Larkin, Roberto Gomperts, Michael Wolfe, and Brent Leback of the NVIDIA® Corporation for general support and discussions, as well as Prof. T. Daniel Crawford of Virginia Tech, Dr. Filippo Lipparini of Johannes Gutenberg-Universität Mainz, and Dr. Radovan Bast of UiT The Arctic University of Norway for providing fruitful comments to the manuscript.

Notes and References

(1) See, for instance, the talk by Jeff Larkin on performance portability from the SC15 conference: `http://images.nvidia.com/events/sc15/pdfs/SC15-Performance-Portability-OpenACC-Larkin.pdf` (accessed September 27, 2016).

(2) Wilson, G.; Aruliah, D. A.; Brown, C. T.; Hong, N. P. C.; Davis, M.; Guy, R. T.; Haddock, S. H. D.; Huff, K. D.; Mitchell, I. M.; Plumbley, M. D.; Waugh, B.; White, E. P.; Wilson, P. PLoS Biol. **2014**, 12, e1001745.

(3) `http://openmp.org/` (accessed September 27, 2016).

(4) See, for instance, the talk by James Beyer and Jeff Larkin on OpenMP vs. OpenACC from the GTC2016 conference: `http://on-demand.gputechconf.com/gtc/2016/presentation/s6410-jeff-larkin-beyer-comparing-open-acc-openmp.pdf` (accessed September 27, 2016).

(5) `http://www.openacc.org/` (accessed September 27, 2016).

(6) Møller, C.; Plesset, M. S. Phys. Rev. **1934**, 46, 618.

(7) (a) Čížek, J. J. Chem. Phys. **1966**, 45, 4256; (b) Čížek, J. Adv. Chem. Phys. **1969**, 14, 35; (c) Paldus, J.; Čížek, J.; Shavitt, I. Phys. Rev. A **1972**, 5, 50.

(8) Purvis, G. D.; Bartlett, R. J. J. Chem. Phys. **1982**, 76, 1910.

(9) Raghavachari, K.; Trucks, G. W.; Pople, J. A.; Head-Gordon, M. Chem. Phys. Lett. **1989**, 157, 479.

(10) Stanton, J. F. Chem. Phys. Lett. **1997**, 281, 130.

(11) (a) Ufimtsev, I. S.; Martínez, T. J. J. Chem. Theory Comp. **2008**, 4, 222; (b) Luehr, N.; Ufimtsev, I. S.; Martínez, T. J. J. Chem. Theory Comp. **2011**, 7, 949.

(12) Asadchev, A.; Allada, V.; Felder, J.; Bode, B. M.; Gordon, M. S.; Windus, T. L. J. Chem. Theory Comp. **2010**, 6, 696.

(13) Wilkinson, K. A.; Sherwood, P.; Guest, M. F.; Naidoo, K. J. J. Comp. Chem. **2011**, 32, 2313.

(14) (a) Miao, Y.; Merz Jr., K. M. J. Chem. Theory Comp. **2013**, 9, 965; (b) Miao, Y.; Merz Jr., K. M. J. Chem. Theory Comp. **2015**, 11, 1449.

(15) Toivanen, E.; Losilla, S. A.; Sundholm, D. Phys. Chem. Chem. Phys. **2015**, 17, 31480.

(16) Yasuda, K. J. Chem. Theory Comp. **2008**, 4, 1230.

(17) (a) Ufimtsev, I. S.; Martínez, T. J. J. Chem. Theory Comp. **2009**, 5, 1004; (b) Ufimtsev, I. S.; Martínez, T. J. J. Chem. Theory Comp. **2009**, 5, 2619.

(18) Asadchev, A.; Gordon, M. S. J. Chem. Theory Comp. **2012**, 8, 4166.

(19) (a) Hohenstein, E. G.; Luehr, N.; Ufimtsev, I. S.; Martínez, T. J. J. Chem. Phys. **2015**, 142, 224103; (b) Hohenstein, E. G.; Bouduban, M. E. F.; Song, C.; Luehr, N.; Ufimtsev, I. S.; Martínez, T. J. J. Chem. Phys. **2015**, 143, 014111; (c) Snyder Jr., J. W.; Curchod, B. F. E.; Martínez, T. J. J. Phys. Chem. Lett. **2016**, 7, 2444.

(20) Yoshikawa, T.; Nakai, H. J. Comp. Chem. **2015**, 36, 164.

(21) Andrade, X.; Aspuru-Guzik, A. J. Chem. Theory Comp. **2013**, 9, 4360.

(22) Cawkwell, M. J.; Wood, M. A.; Niklasson, A. M. N.; Mniszewski, S. M. J. Chem. Theory Comp. **2014**, 10, 5391.

(23) Losilla, S. A.; Watson, M. A.; Aspuru-Guzik, A.; Sundholm, D. J. Chem. Theory Comp. **2015**, 11, 2053.

(24) Kussmann, J.; Ochsenfeld, C. J. Chem. Theory Comp. **2015**, 11, 918.

(25) Liu, F.; Luehr, N.; Kulik, H. J.; Martínez, T. J. J. Chem. Theory Comp. **2015**, 11, 3131.

(26) Nitsche, M. A.; Ferreria, M.; Mocskos, E. E.; Lebrero, M. C. G. J. Chem. Theory Comp. **2014**, 10, 959.

(27) (a) Götz, A. W.; Williamson, M. J.; Xu, D.; Poole, D.; Le Grand, S.; Walker, R. C. J. Chem. Theory Comp. **2012**, 8, 1542; (b) Salomon-Ferrer, R.; Götz, A. W.; Poole, D.; Le Grand, S.; Walker, R. C. J. Chem. Theory Comp. **2013**, 9, 3878.

(28) Wu, X.; Koslowski, A.; Thiel, W. J. Chem. Theory Comp. **2012**, 8, 2272.

(29) (a) Vogt, L.; Olivares-Amaya, R.; Kermes, S.; Shao, Y.; Amador-Bedolla, C.; Aspuru-Guzik, A. J. Phys. Chem. A **2008**, 112, 2049; (b) Olivares-Amaya, R.; Watson, M. A.; Edgar, R. G.; Vogt, L.; Shao, Y.; Aspuru-Guzik, A. J. Chem. Theory Comp. **2010**, 6, 135; (c) Watson, M. A.; Olivares-Amaya, R.; Edgar, R. G.; Aspuru-Guzik, A. Comput. Sci. Eng. **2010**, 12, 40.

(30) Katouda, M.; Naruse, A.; Hirano, Y.; Nakajima, T. J. Comp. Chem. **2016**,

(31) Doran, A. E.; Hirata, S. J. Chem. Theory Comp. **2016**,

(32) DePrince III, A. E.; Hammond, J. R. J. Chem. Theory Comp. **2011**, 7, 1287.

(33) DePrince III, A. E.; Kennedy, M. R.; Sumpter, B. G.; Sherrill, C. D. Mol. Phys. **2014**, 112, 844.

(34) Asadchev, A.; Gordon, M. S. J. Chem. Theory Comp. **2013**, 9, 3385.

(35) Leang, S. S.; Rendell, A. P.; Gordon, M. S. J. Chem. Theory Comp. **2014**, 10, 908.

(36) Ma, W.; Krishnamoorthy, S.; Villa, O.; Kowalski, K.; Agrawal, G. Cluster Comput. **2013**, 16, 131.

(37) (a) Ma, W.; Krishnamoorthy, S.; Villa, O.; Kowalski, K. J. Chem. Theory Comp. **2011**, 7, 1316; (b) Bhaskaran-Nair, K.; Ma, W.; Krishnamoorthy, S.; Villa, O.; van Dam, H. J. J.; Aprà, E.; Kowalski, K. J. Chem. Theory Comp. **2013**, 9, 1949.

(38) Fales, B. S.; Levine, B. G. J. Chem. Theory Comp. **2015**, 11, 4708.

(39) Ufimtsev, I. S.; Martínez, T. J. Comput. Sci. Eng. **2008**, 10, 26.

(40) Valiev, M.; Bylaska, E. J.; Govind, N.; Kowalski, K.; Straatsma, T. P.; van Dam, H. J. J.; Wang, D.; Nieplocha, J.; Aprà, E.; Windus, T. L.; de Jong, W. A. Comput. Phys. Commun. **2010**, 181, 1477.

(41) Pratx, G.; Xing, L. Med. Phys. **2011**, 38, 2685.

(42) Jacob, C. R. J. Phys. Chem. Lett. **2016**, 7, 351.

(43) Song, C.; Wang, L.-P.; Martínez, T. J. J. Chem. Theory Comp. **2016**, 12, 92.

(44) Titov, A.; Ufimtsev, I. S.; Luehr, N.; Martínez, T. J. J. Chem. Theory Comp. **2013**, 9, 213.

(45) Shavitt, I.; Bartlett, R. J. Many-Body Methods in Chemistry and Physics: Many-Body Perturbation Theory and Coupled-Cluster Theory; Cambridge University Press: Cambridge, UK, 2009.

(46) Helgaker, T.; Jørgensen, P.; Olsen, J. Molecular Electronic-Structure Theory, 1st ed.; Wiley & Sons, Ltd.: West Sussex, UK, 2000.

(47) Whitten, J. L. J. Chem. Phys. **1973**, 58, 4496.

(48) Feyereisen, M. W.; Fitzgerald, G.; Komornicki, A. Chem. Phys. Lett. **1993**, 208, 359.

(49) Vahtras, O.; Almlöf, J.; Feyereisen, M. W. Chem. Phys. Lett. **1993**, 213, 514.

(50) Weigend, F.; Häser, M. Theor. Chem. Acc. **1997**, 97, 331.

(51) Almlöf, J. Chem. Phys. Lett. **1991**, 181, 319.

(52) Ayala, P. Y.; Scuseria, G. E. J. Chem. Phys. **1999**, 110, 3660.

(53) Werner, H.-J.; Manby, F. R.; Knowles, P. J. J. Chem. Phys. **2003**, 118, 8149.

(54) (a) Doser, B.; Lambrecht, D. S.; Kussmann, J.; Ochsenfeld, C. J. Chem. Phys. **2009**, 130, 064107; (b) Doser, B.; Zienau, J.; Clin, L.; Lambrecht, D. S.; Ochsenfeld, C. Z. Phys. Chem. **2010**, 224, 397.

(55) Baudin, P.; Ettenhuber, P.; Reine, S.; Kristensen, K.; Kjærgaard, T. J. Chem. Phys. **2016**, 144, 054102.

(56) (a) Schütz, M.; Werner, H.-J. Chem. Phys. Lett. **2000**, 318, 370; (b) Schütz, M. J. Chem. Phys. **2000**, 113, 9986; (c) Werner, H.-J.; Schütz, M. J. Chem. Phys. **2011**, 135, 144116.

(57) Friedrich, J.; Dolg, M. J. Chem. Theory Comput. **2009**, 5, 287.

(58) Kobayashi, M.; Nakai, H. J. Chem. Phys. **2009**, 131, 114108.

(59) Li, W.; Piecuch, P.; Gour, J. R.; Li, S. J. Chem. Phys. **2009**, 131, 114109.

(60) (a) Rolik, Z.; Kállay, M. J. Chem. Phys. **2011**, 135, 104111; (b) Rolik, Z.; Szegedy, L.; Ladjánszki, I.; Ladóczki, B.; Kállay, M. J. Chem. Phys. **2013**, 139, 094105.

(61) Schütz, M.; Yang, J.; Chan, G. K.-L.; Manby, F. R.; Werner, H.-J. J. Chem. Phys. **2013**, 138, 054109.

(62) Riplinger, C.; Sandhoefer, B.; Hansen, A.; Neese, F. J. Chem. Phys. **2013**, 139, 134101.

(63) Eriksen, J. J.; Baudin, P.; Ettenhuber, P.; Kristensen, K.; Kjærgaard, T.; Jørgensen, P. J. Chem. Theory Comput. **2015**, 11, 2984.

(64) Piecuch, P.; Kucharski, S. A.; Kowalski, K.; Musiał, M. Comp. Phys. Comm. **2002**, 149, 71.

(65) (a) Lawson, C. L.; Hanson, R. J.; Kincaid, D.; Krogh, F. T. ACM Trans. Math. Soft. **1979**, 5, 308; (b) Dongarra, J. J.; DuCroz, J.; Hammarling, S.; Hanson, R. ACM Trans. Math. Soft.

1988, 14, 1; (c) Dongarra, J. J.; Duff, I.; DuCroz, J.; Hammarling, S. ACM Trans. Math. Soft. **1990**, 16, 1.

(66) Whaley, R. C.; Petitet, A.; Dongarra, J. J. Parallel Comput. **2001**, 27, 3.

(67) Goto, K.; Van De Geijn, R. ACM Trans. Math. Soft. **2008**, 35, 1.

(68) (a) Tomov, S.; Dongarra, J. J.; Baboulin, M. Parallel Comput. **2010**, 36, 232; (b) Nath, R.; Tomov, S.; Dongarra, J. J. Int. J. High Perform. Comput. Appl. **2010**, 24, 511.

(69) `http://docs.nvidia.com/cuda/cublas` (accessed September 27, 2016).

(70) Lee, V. W.; Kim, C.; Chhugani, J.; Deisher, M.; Kim, D.; Nguyen, A. D.; Satish, N.; Smelyanskiy, M.; Chennupaty, S.; Hammarlund, P.; Singhal, R.; Dubey, P. Debunking the 100X GPU vs. CPU Myth: An Evaluation of Throughput Computing on CPU and GPU. 2010.

(71) Bernholdt, D. E.; Harrison, R. J. Chem. Phys. Lett. **1996**, 250, 477.

(72) Katouda, M.; Nagase, S. Int. J. Quantum Chem. **2009**, 109, 2121.

(73) Katouda, M.; Nakajima, T. J. Chem. Theory Comput. **2013**, 9, 5373.

(74) Maschio, L. J. Chem. Theory Comput. **2011**, 7, 2818.

(75) Werner, H.-J.; Knowles, P. J.; Knizia, G.; Manby, F.; Schutz, M. WIREs Comput. Mol. Sci. **2012**, 2, 242.

(76) (a) Lotrich, V.; Flocke, N.; Ponton, M.; Yau, A. D.; Perera, A.; Deumens, E.; Bartlett, R. J. J. Chem. Phys. **2008**, 128, 15; (b) Deumens, E.; Lotrich, V.; Perera, A.; Ponton, M. J.; Sanders, B. A.; Bartlett, R. J. WIREs Comput. Mol. Sci. **2012**, 1, 895.

(77) Prochnow, E.; Harding, M. E.; Gauss, J. J. Chem. Theory Comput. **2010**, 6, 2339.

(78) (a) Janowski, T.; Pulay, P. J. Chem. Theory Comput. **2008**, 4, 1585; (b) Baker, J.; Janowski, T.; Wolinski, K.; Pulay, P. WIREs Comput. Mol. Sci. **2012**, 2, 63.

(79) (a) Bentz, J. L.; Olson, R. M.; Gordon, M. S.; Schmidt, M. W.; Kendall, R. A. Comp. Phys. Comm. **2007**, 176, 589; (b) Olson, R. M.; Bentz, J. L.; Kendall, R. A.; Schmidt, M. W.; Gordon, M. S. J. Chem. Theory Comput. **2007**, 3, 1312.

(80) (a) Kobayashi, R.; Rendell, A. P. Chem. Phys. Lett. **1997**, 265, 1; (b) Valiev, M.; Bylaska, E. J.; Govind, N.; Kowalski, K.; Straatsma, T. P.; van Dam, H. J. J.; Wang, D.; Nieplocha, J.; Aprà, E.; Windus, T. L.; de Jong, W. A. Comput. Phys. Commun. **2010**, 181, 1477.

(81) Aprà, E.; Harrison, R. J.; Shelton, W. A.; Tipparaju, V.; Vázquez-Mayagoitia, A. Journal of Physics: Conference Series **2009**, 180, 012027.

(82) Olivares-Amaya, R.; Jinich, A.; Watson, M. A.; Aspuru-Guzik, A. In Electronic Structure Calculations on Graphics Processing Units: From Quantum Chemistry to Condensed Matter Physics; Walker, R. C., Götz, A. W., Eds.; John Wiley & Sons, Ltd.: Chichester, UK, 2016.

(83) Ma, W.; Bhaskaran-Nair, K.; Villa, O.; Aprà, E.; Tumeo, A.; Krishnamoorthy, S.; Kowalski, K. In Electronic Structure Calculations on Graphics Processing Units: From Quantum Chemistry to Condensed Matter Physics; Walker, R. C., Götz, A. W., Eds.; John Wiley & Sons, Ltd.: Chichester, UK, 2016.

(84) Rendell, A. P.; Lee, T. J.; Komornicki, A. Chem. Phys. Lett. **1991**, 178, 462.

(85) Rendell, A. P.; Lee, T. J.; Komornicki, A.; Wilson, S. Theor. Chim. Acc. **1992**, 84, 271.

(86) Eriksen, J. J. In Parallel Programming with OpenACC; Farber, R., Ed.; Morgan Kaufmann: Burlington, USA, 2016.

(87) OpenACC-accelerated coupled cluster codes:

https://gitlab.com/januseriksen/openacc_cc_codes (accessed September 27, 2016).

(88) Vysotskiy, V. P.; Cederbaum, L. S. J. Chem. Theory Comp. **2011**, 7, 320.

(89) Knizia, G.; Li, W.; Simon, S.; Werner, H.-J. J. Chem. Theory Comp. **2011**, 7, 2387.

(90) Dunning Jr., T. H. J. Chem. Phys. **1989**, 90, 1007.

(91) Weigend, F.; Köhn, A.; Hättig, C. J. Chem. Phys. **2002**, 116, 3175.

(92) Hättig, C. Phys. Chem. Chem. Phys. **2005**, 7, 59.

(93) Due to a limited amount of memory on the host node (128 GB), the accelerated CCSD(T) calculation on the [ala]-6/cc-pVDZ problem only spawned a total of 14 OpenMP threads instead of the default 20 threads. Out of these fourteen threads, six were devoted to the six K40s, while the remaining eight threads did computations on the host node. Finally, due to timelimit restrictions, the mixed precision CPU-only CCSD(T)/[ala]-6 calculation (using 14 OpenMP threads) was not performed, as were none of the corresponding [ala]-6 calculations in double precision since the memory requirements for these exceeded the size of the host node memory (i.e., the calculations required more than 128 GB).

(94) Appel, A. W.; Ginsburg, M. Modern Compiler Implementation in C; Cambridge University Press: Cambridge, UK, 2004.

(95) Amdahl, G. AFIPS Conf. Proc. **1967**, 30, 483.

(96) (a) Kucharski, S. A.; Bartlett, R. J. J. Chem. Phys. **1998**, 108, 5243; (b) Kucharski, S. A.; Bartlett, R. J. J. Chem. Phys. **1998**, 108, 9221.

(97) Crawford, T. D.; Stanton, J. F. Int. J. Quant. Chem. **1998**, 70, 601.

(98) (a) Kowalski, K.; Piecuch, P. J. Chem. Phys. **2000**, 113, 18; (b) Piecuch, P.; Wloch, M. J. Chem. Phys. **2005**, 123, 224105; (c) Piecuch, P.; Wloch, M.; Gour, J. R.; Kinal, A. Chem. Phys. Lett. **2006**, 418, 467.

(99) (a) Kállay, M.; Gauss, J. J. Chem. Phys. **2005**, 123, 214105; (b) Kállay, M.; Gauss, J. J. Chem. Phys. **2008**, 129, 144101.

(100) (a) Eriksen, J. J.; Kristensen, K.; Kjærgaard, T.; Jørgensen, P.; Gauss, J. J. Chem. Phys. **2014**, 140, 064108; (b) Eriksen, J. J.; Jørgensen, P.; Olsen, J.; Gauss, J. J. Chem. Phys. **2014**, 140, 174114; (c) Eriksen, J. J.; Jørgensen, P.; Gauss, J. J. Chem. Phys. **2015**, 142, 014102; (d) Eriksen, J. J.; Matthews, D. A.; Jørgensen, P.; Gauss, J. J. Chem. Phys. **2015**, 143, 041101; (e) Kristensen, K.; Eriksen, J. J.; Matthews, D. A.; Olsen, J.; Jørgensen, P. J. Chem. Phys. **2016**, 144, 064103; (f) Eriksen, J. J.; Matthews, D. A.; Jørgensen, P.; Gauss, J. J. Chem. Phys. **2016**, 144, 194102; (g) Eriksen, J. J.; Matthews, D. A.; Jørgensen, P.; Gauss, J. J. Chem. Phys. **2016**, 144, 194103.

(101) Kobayashi, M.; Imamura, Y.; Nakai, H. J. Chem. Phys. **2007**, 127, 074103.

(102) (a) Mochizuki, Y.; Yamashita, K.; Murase, T.; Nakano, T.; Fukuzawa, K.; Takematsu, K.; Watanabe, H.; Tanaka, S. Chem. Phys. Lett. **2008**, 457, 396; (b) Ishikawa, T.; Kuwata, K. Chem. Phys. Lett. **2009**, 474, 195.

(103) (a) Guo, Y.; Li, W.; Li, S. J. Phys. Chem. A **2014**, 118, 8996; (b) Guo, Y.; Li, W.; Yuan, D.; Li, S. Sci. China Chem. **2014**, 57, 1393.

(104) (a) Friedrich, J. J. Chem. Phys. **2007**, 126, 154110; (b) Friedrich, J.; Dolg, M. J. Chem. Theory Comput. **2009**, 5, 287.

(105) (a) Ziółkowski, M.; Jansík, B.; Kjærgaard, T.; Jørgensen, P. J. Chem. Phys. **2010**, 133, 014107; (b) Kristensen, K.; Ziółkowski, M.; Jansík, B.; Kjærgaard, T.; Jørgensen, P. J. Chem. Theory Comput. **2011**, 7, 1677; (c) Ettenhuber, P.; Baudin, P.; Kjærgaard, T.; Jørgensen, P.; Kristensen, K. J. Chem. Phys. **2016**, 144, 164116.